The Challenge of Definitively Proving Creativity in AI
https://mindmatters.ai/podcast/ep363

Robert J. Marks:

Greetings and welcome to Mind Matters News. I am your host, Robert J. Marks. We have been talking to Dr. George Mappouras about his very interesting paper called, "The Turing Test 2.O, The General Intelligence Threshold," and it's a new measure for testing whether or not artificial intelligence is truly intelligent.

Now, there've been a number of different for measuring this, classically is the Turing Test. Dr. Mappouras, who has allowed me to refer to him as George, and he to me as Bob, has proposed something called the Turing Test 2.0. And last time we talked about Collays or Collets, I forgot, George. Is it Collay or Collet?

George Mappouras:

Yeah, I'm not sure. I call him Collet.

Robert J. Marks:

Collet.

George Mappouras:

Yeah.

Robert J. Marks:

Okay.

George Mappouras:

I call it this.

Robert J. Marks:

I got both of them wrong, of the Collet. So there's four different measures, and also the Lovelace Test. I have been a big proponent of the Lovelace Test proposed by Selmer Bringsjord of Rensselaer Polytechnic who says that, "Computers will be creative if and when their output is beyond the explanation or intent of the original programmer."

Now this is becoming more and more difficult, because you might get an output for your large language models that you didn't understand and you say, "Wow, this is creative and was not in my intent." But then in order to actually claim that, you would have to go back and look at all the corpus of material that was used to train that large language models and make sure that it was not in that corpus.

So, we've talked about George's model on AI and intelligence, and do you want to contrast that, George, with the Lovelace Test?

George Mappouras:

Yeah, so one thing we talked in the previous episodes, right? Was this what you refer to as a strike of genius, right?

Robert J. Marks:

Yes.

George Mappouras:

And Lovelace Test come very close to the same definition as the one I give, but I think there's some fundamental differences. And the first one, I think, we have to give the chance to AI to pass the test. So, although very similar definitions, what I don't like about the Lovelace Test, or what I would disagree with, let's say, is that let's say I do design an AI that is truly intelligence, right?

So that's my, or that's at least what I claim, then I cannot be surprised by its output. So it's kind of a contradictive term as in, how can I be surprised if that's my intent?

Robert J. Marks:

Well, I would maintain George, that surprise is not the same thing as creativity. I've written a lot of programs and I'm kind of surprised at the output, but I can usually go back and explain why that surprise happened. It's something that I programmed it to do. I might ask it to search through a billion different possibilities for a solution, and it comes out with one and said, "Wow, I wouldn't expect that, I'm surprising."

But then I go back and look at the details and I find that it's right. So, I don't know if I would equate surprise with, two, creativity. So go ahead. That's my two cents.

George Mappouras:

Yeah, no, I agree with that. What I'm saying though, in the Lovelace Test, if it's truly creative, I won't be able to explain how it came to that result.

Robert J. Marks:

Right.

George Mappouras:

But if I designed it like that, then I wouldn't be able to explain it anyway. So, you see what mean? So, even if I intentionally it to be creative, right? Then if there is such a path, then I'll be able to say, "Here," because I design it like that. So I feel like that's where the contradiction comes.

But moreover is the problem is what does it mean not be able to explain. So for example, in the previous episodes, we gave the example about the anecdote of the Newton, and Newton observing the apple falling. And as it reaches the ground, right, and then from that extracts new knowledge that objects attract each other.

So how much information can I give until I say, "Oh yeah, I can see how it came with this result. You looked at this information and it extracted some pattern," right? Because it's a problem you said before, I have to go back and check if that information was already in the data set. So how much information am I allowed to give you?

So, this is not well-defined in the Lovelace Test. Where in my proposal, it's defined, because I define what that means. I mean, that means new functionality. So I can check what functions a system can do. I can find one you cannot, and then I can check, do you have information that you can extract this new functionality out of?

So for example, we talked about the 6:30 example, the clock that points at 6:30 where the AI can describe you precisely in words, but it cannot generate an image because it cannot extract the knowledge from the information it actually provides to you. So, describes where the hands of the clock should be, but it cannot extract the functionality.

And it is not like I disagree with the Lovelace Test, as in I feel like things are not well-defined. Right? So that's why I felt the urge to be like, we really have to define precisely what does it mean, "creativity"? What does it mean, "general intelligence." So-

Robert J. Marks:

So let me-

George Mappouras:

... that's where-

Robert J. Marks:

... ask you, you used the phrase, suppose that we wrote software that was designed to be creative.

George Mappouras:

Yeah.

Robert J. Marks:

Is that possible? How would we do that? How we would design an AI to be creative? I have no idea how to do that. Is it possible?

George Mappouras:

I agree that I don't see a path to it.

Robert J. Marks:

Okay.

George Mappouras:

But I want to leave this open as in, I don't want to be too restrictive. Because then the counter argument will be your definition is such as to make it impossible. So I want to avoid this step, which sometimes people fall into, right, we fall into because we have our own preconceived biases.

For example, I don't see a path to AGI, but I want to leave an open door where yes, if it's there, I will be able to detect it. It'd be like, "Okay, that as well."

Robert J. Marks:

You'll be able to detect it.

George Mappouras:

Yeah.

Robert J. Marks:

Okay. Okay. That's fair enough. What about application to other things? We actually talked about the idea of flash of genius-

George Mappouras:

Right.

Robert J. Marks:

And I think that in order for AI to be creative, that's indeed what we're looking for in all of these tests for intelligence, we're looking for a flash of genius. And I think that that applies to human beings. Roger Penrose, in his book, The Emperor's New Mind, points out a number of different places where this flash of genius has occurred. You and I think anybody creative has had this flash of genius where you actually come up with a solution and it was nothing like you thought about and you go, "Oh my gosh, where did that come from?"

And so the flash of geniuses do indeed occur, and so we can have this. And so these different models of intelligence, I think it's important that we apply it to other things. Now, you mentioned that your Turing Test 2.0 can be applied to other things. Could you unpack that a little bit?

George Mappouras:

Yeah, so theoretically at least, right, that's why let's say I really like that specific definition of functionality, like emerging functionality, let's say, generating functionality. And because already clearly it can be applied to any system. And what I mean by that, for example, one way I came with this idea is by saying, we talked about in the previous episodes, we look at humans as a system, and we see how their technology evolved. Right? So they advanced.

We can look at other animals, and at least with the data we have today, we don't see them. We see that other animals live as they always lived. An example would be you might have a pet dog, right? And you learn it, you teach it tricks. And it's learning the tricks, but then it won't learn anything else. Right?

It will still live in the same way. It won't take these tricks and try to build upon them. A similar example is you might have seen videos where people teach the sign language to apes, and try to communicate with them. Okay, very nice, very interesting. It shows some aspects of intelligence, but is it general intelligence? If I take this ape and take it back to nature, will he use this communication to build a society? Will there be an advancement, like, we will see them to use new technology, right? Do some innovative things, new functionality. That's why functionality is important. And it's the same thing.

And if the same thing comes to anything else, like for example, one thing we mentioned is that humans typically have different levels of intelligence in different aspects like art-

Robert J. Marks:

Sure.

George Mappouras:

... sports, science, and you can see that, for example, I like the example of NBA, like if we look at NBA today, how it's played the basketball game, right?

Robert J. Marks:

Oh, NBA. Okay, yes.

George Mappouras:

Yeah, NBA, it's very different. If you look at the '60s and what happened. Right? So, people started rethinking the game and saying, "Wait a second, maybe I don't need as much physical contact. Maybe I can shoot from the three point line or further back. And statistically, even if I miss more because it's more points, it gives me a bigger chance to win."

And people arrive to new functionality, new play styles that were not there before. We see athletes doing new moves to trick their opponents. Right? This is new functionality and the game evolves. And you can see, and if you look at a game back then and now, very different. Same with music, same with let's say, paintings. We can see new trends coming up, and I think this is exactly the new functionality they do.

Sometimes people might think that we see that with AI, but what we see in AI is that it's a combination of what already exists right now. That's why I propose in the Turing Test 2.0, I propose specific rules that you have to apply, like having some functionality, having some information that you haven't extracted functionality out of it, and then no new information coming to the system.

Can it extract new functionality out of it? So for example, you could apply that let's say by saying, in going to the painting world and say, "What was the different types, the different methods people used to paint?" Right? So make, paint pictures. And maybe it would go back to, I don't know, the 800, and we only train the model with these type of images.

Or you can go even further back, let's say the Stone Age, and you only have the pictures on the walls that the primitive humans were doing. And you say, "Okay, I'll only give you these images." And maybe all the information that the humans had at that time. Can you come with more innovative images?

Robert J. Marks:

Yes.

George Mappouras:

So this is the idea of new functionality. You can observe it everywhere. You have some type of intelligence.

Robert J. Marks:

I think part of the intelligence you're talking about is the creative dimension.

George Mappouras:

Exactly.

Robert J. Marks:

I think that animals probably have a degree of creativity. You hear about the people that put a monkey in a little room and they hang bananas and they put a stick there, and the monkey figures out how to use the stick to knock the bananas down. That's probably a low level of creativity, but I think the higher level of creativity is probably limited to human beings and possibly our creator.

What blows my mind is that everything that we have now from computers to cars and this great lifestyle that we live is made with stuff that was here on the earth 3,000 years ago.

George Mappouras:

Exactly.

Robert J. Marks:

And oh my gosh, we have taken that. We have used our creativity to form the society that we have now. And this takes human creativity and this human creativity is just awesome at what it can do. And now we're seeing this application to computers and artificial intelligence. It's just astonished, all of the silicon, man, it was lying around on beaches, I think. Right?

George Mappouras:

Exactly.

Robert J. Marks:

And it was all there and we just figured out how to use it. And we made cars and computers and houses and skyscrapers and boats and explosives, and all sorts of things.

George Mappouras:

Yeah. Good and bad.

Robert J. Marks:

What's that?

George Mappouras:

Good and bad applications.

Robert J. Marks:

Good and bad. Exactly. I mean it was there, but I do believe it does have its sort of limitations.

George Mappouras:

I like what you said. We learned how to use it. Right? It was always there. The silicon always had the same purposes.

Robert J. Marks:

It's always there. That's right.

George Mappouras:

And we learned what that means. How, what is the functionality we can extract out of it?

Robert J. Marks:

Yes.

George Mappouras:

And that's the point.

Robert J. Marks:

Yeah. And that goes back, the simple example you gave about Newton, seeing the apple fall from the tree, that old myth. And now he was able to, how he was to use creativity to do great things. I don't think computers will ever be able to do that. And at least I would agree with you. I think you state it more safely. You don't see a path to it. I definitely do not see a path to it.

So let me bring up a topic that I didn't put down in the notes for you, but I'd be interested in your take on it. There was this recent article, I believe it was in Forbes, and I'll try to find it and put it in the podcast notes that said, "We shouldn't be studying large language models, we should be studying puppy dogs."

And they say, "Why should we be studying puppy dogs?" "Well, we should be using brains." It does turn out that there's something that comes from the brain, and I believe, in fact, I edited a book called, or co-edited a book called Minding the Brain, that I believe that the mind is more than the brain, that we're more than computers made out of meat.

The inability of AI to create, I think, is evidence of that. It isn't a proof, but I think it's certainly evidence. And so they were making the point that maybe in the future we could actually grow brains. And right now they take pigs and they grow like pancreases and livers and stuff. I don't know the details, but I ask a guy, "Could you grow a human brain on a pig?"

And they says, "Yeah, we could." And I just wonder what we could do with a computer, with a human brain made out of meat, and whether that would be creative or not. I don't think theologically or philosophically it would have a soul or a spirit at all. But is there something that we could tap in there? Do you have any thoughts on that? Have you ever thought about that?

George Mappouras:

Yeah, so lots of thoughts. Very interesting subject. I think when we think it, I like to break it down. And that's one aspect again of the Turing Test 2.0 is let's break it down. It means, what does it mean to be intelligent, we set this new functionality. Okay, there's two paths to it.

Either this thing is actually able to be done algorithmically or not.

Robert J. Marks:

Yes.

George Mappouras:

There's only two possible outcomes. So if you break it down like that, that's why I like to define it precisely so I can break it down like that. And then it's like, if it can't be done algorithmically, then eventually we will do it. It doesn't matter of time. You see what I mean? If it's an algorithm, maybe not now, maybe a hundred years from now, it doesn't matter. It will come down. We can't even randomly just change things until we figure it out.

Robert J. Marks:

Maybe. Okay.

George Mappouras:

So it's the same thing with the brain. If it's an algorithm, then the brain itself should be sufficient. And if you put it in any other system, it should be sufficient. So if it's not, so if we go to the other end, then okay, we examined this case. But if it's not, then I think this is fundamentally groundbreaking.

If it's not, because then if algorithms cannot do it, then there's no path to it. No, sometimes we think about algorithms and especially with AI, we think about it as a mysterious thing. What is an algorithm? Right? Like AI, first of all, it's algorithms, because sometimes people forget about that.

Robert J. Marks:

In fact, if I could interject really quick-

George Mappouras:

Yeah.

Robert J. Marks:

... this is a point that Roger Penrose makes in his book, The Emperor's New Mind. He says, "The computers are limited to algorithmic things." And we know, shoot, back in the '30s, Alan Turing showed the halting problem was non-algorithmic, and now we know a bunch of things which are non-algorithmic.

George Mappouras:

Right.

Robert J. Marks:

There are a bunch of problems that computers will never solve. And that's where the human mind comes in, I believe.

George Mappouras:

Exactly. So if we think about the algorithm, what is an algorithm? Somebody Googles a term, you can find a definition. It's not super precise, but you can find a definition. It's a finite steps. Finite steps, a process with finite steps that comes to an outcome. So you put an input, there's finite of processing steps, and the finite is very important.

Because if it's not finite, then you never get an answer back.

Robert J. Marks:

Right.

George Mappouras:

So, if it's finite like that as we said, then all there it is this algorithm and you can map it into different things, different Turing machines. The Turing machine might be a brain, it might be modern computer, it might be anything we can define as a Turing machine. However, if it's not, then that's a big, big issue.

It means there's something else. Some people might disagree on what it is, but yeah, definitely then it is not the algorithm. The issue is not the algorithm, but in order to get either or, that's what I don't like when people talk about these things. In order to get either or, we have to actually first define what we're looking for and then prove it's come to the build.

If we can prove for or against, which is then you have to actually prove it mathematically. Until then, the best thing we have is testing, and seeing how close you're coming to this test. If somebody passes the test, this is proof enough that it exists. If we don't have anybody to pass the test, we are still though in the, we don't know.

Maybe somebody pass in the future, maybe not, yeah, which I think it creates this interesting debate.

Robert J. Marks:

Okay, back to your paper. You propose advanced versions of your test, like the single-discipline test and the generational test. How might these overcome weaknesses in simpler examples?

George Mappouras:

Yeah, so in the last episode, we talked about the test that I show off in the paper, which is they're nice intuitive tests like the clock where they cannot produce a specific time, an image of a clock showing a specific time, but only tend to, because that's all the images they have seen.

Robert J. Marks:

Yeah, yeah.

George Mappouras:

And as you said, people can come and fix these things. We just give you a better data set. Right? We balance the data set, you can fix these things. So these things can go away. So this is one of the problems. The other problem is sometimes it's very hard to figure it out, as you said, what was in that training data set, right? In order to know if this is a stack of genius and new knowledge, or is it something that was already in the data set?

And that's why I propose to other tests. I don't test them, but it's out there if somebody wants to use them. The single discipline test is what we talked about. The idea of it was actually a podcast. I mentioned it before where Sam Altman was in a podcast called Huge Conversations, interviewer asked me-

Robert J. Marks:

By the way, if you could send me the link to that, I'd be happy to put it in the podcast notes for-

George Mappouras:

Yeah, yeah. Of course.

Robert J. Marks:

... people that are interested. Thank you.

George Mappouras:

And the interviewer asked him, "Let's assume I take your new, I think the GPT-5 now I train it only with Newton's time of physics. So whatever Newton knew, right? And Newton had a specific theory about how gravity works, that today we know it's wrong. Einstein told us Newton's theory of gravity is wrong. Here's how it's done.

Will GPT-5 having done knowledge, be able to gradually come to the conclusion? Sam's answer there was pretty much a little bit avoiding the answer, but he said pretty much that sometimes what we need is more data. But that's why I like the Turing Test 2.0, I'll give you the data. I just don't tell you how to interpret it. That's the idea.

So I don't want you to go to Einstein's theory. I'll do you a smaller step. I'll give you, let's say, a more relaxed version, invent something new after Newton, anything, anything that came after that point. So I train you with a Newton's time, physics, only physics. So it's easier to expect results. That's a single-discipline test only physics.

And then I want you to figure out the next thing, whatever is that tiny new invention that was made, and we kind of touched on a bit on the previous episode or if I'm not mistaken, where we talked about images. And you can have, I'll train you on the images of the primitive humans. It's all, let's say images on stone. Right? And then I'll see, can you come with new type of pictures, right?

Realism or something else. The same thing with any other aspect of intelligence like sports. It can be science, it can be medicine, anything. So this is the single-discipline test. I'll go back in history in a specific time where I knew what was all the information, all the knowledge at that time. I give all that knowledge, all the functionality to the system, plus the information we had at that time, which as always in, for example, currently we have some specific amount of knowledge.

We're in the frontier of knowledge, right? There is some-

Robert J. Marks:

No.

George Mappouras:

... of the things we have to remember, some of the things we know today are wrong. There's some functionality we have it wrong.

So if I go back in time, I can do that. Come on the time at that frontier, there's some things that we knew back then that were wrong and we corrected them. Can AI do that? Can get us a tiny leap? And it's easier to observe because we know what were the next steps. So it's easier to see what we're looking for.

But even that test is very hard. What does it mean to have all the knowledge of that time? Sometimes it's very hard to define that. It's very hard to make sure you didn't cheat even accidentally, because you have huge amount of data, even accidentally, you might give data that you're not supposed to and train your AI with data you were not supposed to.

So that's why then it comes a generational test. And a generational test is the same idea as humans. We see that we have a generation of humans teaching the next generation, the younger generation, the younger generation comes up, invents new things, teaches the next, and it keeps going. So this is evident when we look at humans and how the everything, technology, music, art, sports evolved. We invent new things.

So, that is the ultimate, I would call, the ultimate test. If an AI can only using the data it can generate, generated AI data, can train the next generation of AI, and then eventually come with new functionality that wasn't there before being able to solve the problem that wasn't able to solve before, then we're going to have passing the Turing Test 2.0.

And actually the funny part, as long as I thought that I went to search it, I'm like, no way. Somebody must have thought of that. And what it was fun is that people actually use that, not for the purpose of figuring out if a model is intelligent, but rather than to accelerate training. Right?

So that rather than me having to give you data, let AI give you the data. And what they observe is that we have generations of training. Not only the models don't become better, they actually become worse.

Robert J. Marks:

Well, in fact, I think that this phenomenon is, I don't know if you've heard the term "model collapse."

George Mappouras:

Exactly, yeah, yeah.

Robert J. Marks:

Where they have tried to use AI to train better AI, to train better AI, and just, after a few generations it becomes a blubbering idiot. And-

George Mappouras:

It does.

Robert J. Marks:

... so I think it's evidence, I don't think it's proof, but I think it's evidence that AI doesn't have the ability to write better AI. And this goes back to the idea that it can't be creative. So.

George Mappouras:

Exactly, that's exactly my point. And it's good to go to the, let's say, new instance of this thing, because there are solutions out there. For example, people say, "Okay, I'll train you with generative data, but also I'm going to introduce some real data from humans in order to avoid the model collapse." Right?

But this thing, although it solves the problem or mitigates the problem, it breaks the third rule of Turing Test 2.0.

Robert J. Marks:

I see, yes.

George Mappouras:

Because you introduce functionality externally. And I want to mention that I didn't come with these rules in order to counter these arguments, I came with the rules first, and then I saw the test. I was like, "Oh yeah, of course, in order to avoid this problem, you break the rules of the test." And I find it very interesting.

Robert J. Marks:

Very good. Let me end by asking you a question that a lot of people are talking about. There's people that believe, of course, this super intelligence is going to emerge, including Ray Kurzweil of Google, Yuval Noah Harari, the late Stephen Hawking believed in this, but his Nobel Prize buddy, Roger Penrose did not believe in this. What are going to be the societal and ethical implications if one day AI demonstrably crosses your general intelligence threshold? What's going to be the impact? It's kind of scary, isn't it?

George Mappouras:

It's definitely, I want to say weird to even think about it, right?

Robert J. Marks:

Weird, yeah.

George Mappouras:

Because what does that mean? It means, I guess it means we have, first of all what it means, all right, it means that our type of creativity is just an algorithm. Right? A lot of people already believe that. That's why they think. I think one of the things that people believe that we're going to get there because if we're doing it and if we're just an algorithm-

Robert J. Marks:

Well, I tell people, these people believe that we're computers made out of meat, and therefore if we can create, then certainly if we replicate ourselves in silicon that it can create also. And I think that that's a vacuous model.

George Mappouras:

Right? But it is logical. It's a logical extrapolation.

Robert J. Marks:

Well, it's logical if you, depending on your ideology.

George Mappouras:

Yeah, I agree. But what I'm saying is, it's kind of logical extrapolation. And the thing is, if this is really true, right, then you have a machine, we have to think about how humans run, right? Humans use these neurons that need, let's say milliseconds.

What is a millisecond for a computer is ages-

Robert J. Marks:

Yeah, it is.

George Mappouras:

Milliseconds to communicate. What is this? We're talking about nanoseconds, right? The bandwidth of a computer. How much information can you intake the latency from the time I asked you to do something until you give me a response. So if you truly have an algorithm that reaches that, and it's so much superior in hardware because it's not just software, it's also hardware, then this is going to be like a thing we cannot not even imagine.

We might be trying to imagine, but this going to be, it might be even the case where we cannot even understand it anymore. You know what I mean? And it might even be less useful in the end because-

Robert J. Marks:

Less useful?

George Mappouras:

Less useful, because imagine-

Robert J. Marks:

Okay.

George Mappouras:

... if you go to a kid that is in the first grade and you try to teach it general relativity-

Robert J. Marks:

Yes?

George Mappouras:

... what is it going to learn? Nothing. You see me? So it has to be at least gradual. It's going to be our teacher. So it has to be a good teacher. If it's not a good teacher, then it's not going to be good. Very useful. You see me?

Robert J. Marks:

Yep.

George Mappouras:

But yeah, definitely. It is very hard to predict what's going to happen. That's just a take. Yeah. But I think all this talk though, however, as you said, it's up in this assumption. Is this possible? And currently, I don't think we have at least a mathematical proof that this is possible. It's more like a belief or a hunch-

Robert J. Marks:

I would say a faith.

George Mappouras:

Yeah, it's a belief. It's a faith.

Robert J. Marks:

Yeah.

George Mappouras:

Yeah.

Robert J. Marks:

So, yeah, there's this number called Chaitin's number, which is astonishing from algorithmic information theory. If you knew this one number, you could prove or disprove all of the open mathematical problems that can be proved, that can be disproved with the single counterexample. It's an astonishing number.

And Chaitin proved that it exists and it exists. C++ has a Chaitin number, as does Python, but you can also prove that it's unknowable. It might be that there are some things which are unknowable, some things are true because they're true and are beyond computing or breaking down. So it could be that that's another answer to this. I don't know.

George Mappouras:

Yeah.

Robert J. Marks:

Any final thoughts?

George Mappouras:

Yeah, I think what is after this, right? I think this is the nice question right now, right? Okay. If you have a good definition of general intelligence, then I think the next step is what exactly we discussed right now, try to look at it and see, can we map it to an algorithm or not?

Robert J. Marks:

Yes. I believe if you define AGI as climbing that mountain and standing on the top of it, I think that currently the large language models get you to that top of that mountain quicker than anything else. I mean, all of the information that you have in GROK or Chat GPT is available in the US Library of Congress.

It would just take you a long time in order to extract it in order to get to the top of that mountain. But Chat GPT gets you there immediately.

George Mappouras:

Yeah. So one more point that I wanted to add is, because we talked about what is going to be in vision if this thing actually becomes intelligent. One thing we have to always remember is the opposite though. What if it's not? Right? Because if it's not, the problem is, the other problem with AI, typically people think about AI, they're like, "Oh, the doom that is going to come, if it gets super intelligent, it will destroy us all."

But we don't think about the other side of the coin. If it's not truly intelligent, there's a fear of a false prophet. And what do I mean by that is we have this in our society multiple times. For example, we had the anchorman, the newsman, whatever the newsman says is true because look at him, he's on the TV dressed nicely. He knows what he's talking about.

Then we had scientists say, and oh yeah, if scientists say it, I don't know if you know the nice wordplay that says, "Scientists say that people are more likely to believe something if scientists-

Robert J. Marks:

Yeah.

George Mappouras:

... say," right? And then after the scientists say, maybe we have, I Googled it, and then now we have AI setting. So this false prophet is an actual problem, especially when people over-exaggerate its abilities. And then we might have this problem where we keep relying on AI that gives us data we already know.

And then we keep putting out data that is generated by AI. AI uses this data to train again, because it's on the Internet, the more data generated by the more the data that AI uses to train is artificially generated.

Robert J. Marks:

Yeah.

George Mappouras:

Then we have this closed loop where no new information, no new knowledge is produced. And then we have the same problem of a stagnation of information stagnation, of a model collapse, where we don't produce nothing else because we rely on something we assume is truly intelligent, general intelligent, but it's not.

And I think that's the other aspect of it, they'll have to be careful. As long as they don't pass the test, we have to remember, all they produce now is what we know today, the frontier of knowledge. So in other words, if the frontier of knowledge today was Galileo's time, AI would insist that the earth is not moving. It's in the center of the universe, and everything spins around it. You have to remember that.

Because then it becomes this problem where, "No, AI said it." "Yeah, but AI says what we think right now." That's why if you see AI playing chess, it plays like then average person, because that's the average knowledge. Doesn't play like an expert.

Robert J. Marks:

Yep.

George Mappouras:

Right? And I think that's very important to the mentor.

Robert J. Marks:

Wow. Well, thank you, George. I think this has been just a fascinating conversation. We've been talking to Dr. George Mappouras, known to his American friends as George, about his new interesting paper, Turing Test 2.0, The General Intelligence Threshold. Again, we're going to make a link to that on the podcast notes so that you can read that as well as make available to links to other literature.

And if George, if you could send me the URL to that podcast, we can include that too.

George Mappouras:

I will do, yeah.

Robert J. Marks:

And so thank you again. This has been a wonderful time. And until next time on Mind Matters News, be of good cheer.

Announcer:

This has been Mind Matters News with your host, Robert J. Marks. Explore more at mindmatters.ai. That's mindmatters.ai. Mind Matters News is directed and by Austin Egbert. The opinions expressed on this program are solely those of the speakers.